Seeing the Forest through the Trees: Adaptive Local Exploration of Large Graphs

Robert Pienta[◦] Zhiyuan Lin[◦] Minsuk Kahng[◦] Jilles Vreeken[•] Partha P. Talukdar[§]
James Abello[‡] Ganesh Parameswaran[¶] Duen Horng (Polo) Chau[◦]
[◦] Georgia Institute of Technology, [•] Max Planck Institute for Informatics and Saarland University,
[§] Indian Institute of Science, [‡] Rutgers University, [¶] Yahoo! Inc.
{pientars,zlin48,kahng}@gatech.edu, jilles@mpi-inf.mpg.de, ppt@serc.iisc.in,
abello@dimacs.rutgers.edu, {ganeshcs,polo}@gatech.edu

arXiv:1505.06792v2 [cs.IR] 22 Jul 2016

ABSTRACT

Visualization is a powerful paradigm for exploratory data analysis. Visualizing large graphs, however, often results in a meaningless hairball. In this paper, we propose a different approach that helps the user *adaptively* explore large million-node graphs from a *local* perspective. For nodes that the user investigates, we propose to only show the neighbors with the most subjectively interesting neighborhoods. We contribute novel ideas to measure this interestingness in terms of how surprising a neighborhood is given the background distribution, as well as how well it fits the nodes the user chose to explore.

We introduce FACETS, a fast and scalable method for visually exploring large graphs. By implementing our above ideas, it allows users to look into the forest through its trees. Empirical evaluation shows that our method works very well in practice, providing rankings of nodes that match interests of users. Moreover, as it scales linearly, FACETS is suited for the exploration of very large graphs.

Categories and Subject Descriptors

H.2.8 [**Database management**]: Database applications–*Data mining*

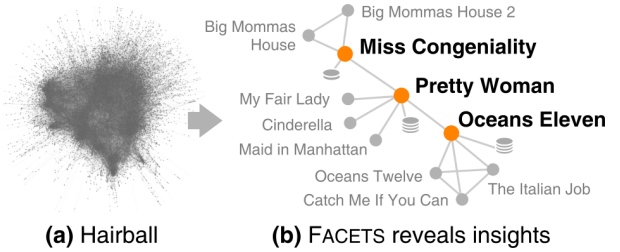
Keywords

Graph Visualization; Data Exploration; Adaptivity; Serendipity

1. INTRODUCTION

Large graphs are ubiquitous. They are natural representations for many domains, and hence we find graph structured data everywhere. As data collection becomes increasingly simple, and many domains remain complex, real-world graphs are growing extremely large and rich with data. These graphs may have thousands or more of attributes and scale up to and beyond millions of nodes and billions of edges. It is fair to say that many graphs are in fact *too big*; exploring such large graphs, where the goal of the user is to gain understanding, is a highly non-trivial task.

Visualization is perhaps the most natural approach to exploratory data analysis. Under the right visualization, finding patterns, deciding what is interesting, what is not, and what to investigate next are all easy tasks – in a sense the answers 'jump to us' as our brains are highly specialized for analyzing complex visual data. It is therefore no surprise that Shneiderman's mantra of "overview, zoom & filter, details-on-demand" [23] has proven to be successful in many domains [13, 22, 23].



(a) Hairball **(b)** FACETS reveals insights

Figure 1: (a) The Rotten Tomatoes movie graph shown using conventional spring layout (an edge connects two movie nodes if some users voted them as similar). Even for this relatively small graph of 17k nodes and 72k edges, a global visualization does not provide much insight. (b) A better way, using our FACETS approach, focuses on movies that are the most subjectively interesting, surprising, or both. For example, FACETS suggests *Pretty Woman* (romantic-comedy) as a interesting, suprising related movie of *Miss Congeniality* (crime-comedy).

Visualizing large graphs in an intuitive and informative manner has proven to be rather difficult. Even using advanced layout techniques plotting a graph typically leads to a useless "hairball" from which nothing can be deduced [13, 14]. This is even the case when we plot graphs with only thousands of nodes (see Figure 1(a) for an example). Instead of plotting the whole graph, visualizing only part of the graph seems more promising [28, 5]. Doing this naively leads to the same problem. Because real world graphs are often scale free (follow a power law degree distribution [8]), even a single hop expansion from a selected node can be visually overwhelming.

We take a different approach. We propose to *adaptively* explore large graphs from a *local* perspective. That is, starting from an initially selected node – e.g., explicitly queried by the user, or proposed by an outlier detection algorithm [2] – we propose to only show the most interesting neighbors. We identify these by their *subjective interestingness*, by how surprising their neighborhood distribution is (e.g., do neighbors' degree distributions follow a power law like when considering all nodes?), as well as by how similar this distribution is to those of the nodes the user already visited during the exploration. By only showing those parts of the graph that are most informative to the user. we keep the view

clean. By being adaptive, FACETS allows every user to explore those facets of the graph that are most interesting to them.

We call our adaptive approach FACETS — our idea is a significant addition to existing works that aim to recommend individual nodes to users; instead, we steer users towards local regions of the graphs that match best with their current browsing interests, helping them better understand and visualize the graph at the same time.

To illustrate how FACETS works in practice, consider our user Susan who is looking for interesting movies to watch (see Figure 1), by exploring a Rotten Tomatoes movie similarity graph with 17k movies. In this graph, an edge connects two movie nodes if users of RottenTomatoes voted them as similar films. Susan has watched *Miss Congeniality*[1], a crime-comedy that stars Sandra Bullock as an FBI agent who thwarts terrorist efforts by going undercover, turning her rude unflattering self into a glamorous beauty queen (see Figure 1b). FACETS simultaneously suggest a few movies that are interesting and surprising to her.

Matching Susan's interest, FACETS suggests the *Big Mommas House* series, which also have undercover plots and are interestingly like *Miss Congeniality*. They both share low critics scores, but high audience scores (i.e., most critics do not like them, but people love them). To Susan's surprise, FACETS suggests *Pretty Woman*, which is quite different (thus surprising) — a romantic-comedy that has both scores from the critics and the audience. But, there is still more subtle similarity (thus still drawing Susan's interest); both films share a Cinderalla-like storyline, which explains why the two movies are connected in the graph: Sandra Bullock goes from a rude agent to a beauty queen; in *Pretty Woman*, Julia Roberts goes from a prostitute to a fair lady. In fact, *Pretty Woman* is a classic, exemplar romantic-comedy; many movies follow similar story lines (e.g., *Maid in Manhattan*). Thus, *Pretty Woman* has very high degree in the graph, unlike *Miss Congeniality* which is a niche genre; this also contributes to *Pretty Woman*'s surprisingness.

Through *Pretty Woman*, FACETS again pleasantly surprises Susan with *Oceans Eleven*, which also stars Julia Roberts, and is in a rather different light hearted crime or heist genre, introducing Susan to other very similar movies like *Oceans Twelve* and *The Italian Job*. Figure 1(b) summarizes Susan's exploration. If Susan were to use a conventional visualization tool to perform the same kind of movie exploration, she would likely be completely overwhelmed with an incomprehensible hairball visualization (as in Figure 1(a)).

The key contributions of our work include:

- A framework for locally exploring a graph without clutter, showing only the most subjectively most interesting nodes, and hence being *adaptive* to the users' interests.
- A formal notion of subjective interestingness for graph exploration taking both divergence between local and global distributions, and similarity to explored nodes into account.
- A measure of surprise over neighborhoods – rather than local node attributes – to draw users in the direction of graph areas with subjectively interesting content.
- A highly scalable method, FACETS, for adaptively exploring very large graphs in a visual environment. Experimental evidence demonstrates the effectiveness of FACETS.

The rest of the paper is organized as follows. First, in Sec. 2 we formalize the problem, introduces our notions of interestingness, and propose our FACETS solution. Then, in Sec. 3 we present our ideas as an integrated approach, with visualization and algorithms working closely together. We empirically evaluate in Sec. 4. We discuss related work in Sec. 5, and end with conclusions in Sec. 6.

[1]One of the most frequently rated movies on Netflix

2. FACETS: ADAPTIVE GRAPH EXPLORATION

In this section, we formalize the problems we aim to solve through our FACETS approach to achieve adaptive exploration. Then, we describe our main ideas, and describe our solutions.

To enhance readability, we have listed the symbols used in this paper in Table 1. The reader may want to return to this table for technical terms meanings used in various contexts of discussion.

Symbol	Description
v_i	Node i
D_{JS}	Jensen-Shannon Divergence
D_{KL}	Kullback-Leibler Divergence
s_i	surprise-score for node v_i
r_i	interest-score for node v_i
$\hat{S}_a$	surprise scores for all neighbors of v_a
$\hat{R}_a$	interest scores for all neighbors of v_a
f_j	a j-th feature for nodes
λ_j	weight of feature f_j
$L_{i,j}$	neighborhood dist. of node v_i for feature f_j
G_j	global dist. for feature f_j
U_j	user profile dist. for feature f_j

Table 1: Symbols & Notation

2.1 Problem Definition

The input is a graph $G = (V, E, A)$ where A is a set of numerical or categorical attributes. Each node $v_i \in V$ has a corresponding attribute value for each attribute (feature) $f_j \in A$ (e.g., degree). Our approach works with both categorical and numerical attributes. We assume there are no self-edges.

We solve the following problem with FACETS:

Definition 1. **Node-Sequence Aware Ranking.** Given a starting node v_a, a sequence of nodes $V_h \subset V$ in which a user has shown interest, how can we find the top-k nodes among the neighbors of v_a that balance (1) similarity by features to the sequence of V_h nodes (subjective interest) and (2) uncommon compared to the global distribution (surprise).

Graph exploration is an interactive and iterative process, where the user incrementally explore larger parts of the graph. FACETS solves the above problem repeatedly.

A common approach to rank nodes is by their *importance* scores, which are often computed using PageRank [20], Personalized PageRank [11] or random walk with restart [27].

However, there are other ways to rank the nodes, like using surprise or interest [19, 2]. We have chosen to rank nodes by their surprise and user-driven interest rather than by the more conventional importance metrics. We chose surprise, because serendipitous results and insight do not always come from the most topologically important nodes [19]. We made FACETS adaptive, because what makes a nodes interesting varies from person to person and so too should be the criteria for ranking. For each node we suggest a combination of the most surprising and most interesting neighbors at each step of the journey.

2.2 Feature Distributions

FACETS uses feature-based surprise and interest in order to guide the graph exploration process. Even when a dataset does not contain node-level features, we can derive node features by using established approaches like PageRank, centrality measures or labels drawn from clustering approaches. This means that even without a set of initial features, it's still possible for FACETS to guide graph exploration.

FACETS requires a compact representation of feature distributions. Histograms are a natural and computationally inexpensive way to represent distributions. Our approach can consider any histogram, regardless of the binning strategy – e.g., equi-width or equi-height binning – used to infer the histogram. Here, we opt to use the parameter-free technique by Kontkanen and Myllymaki [15] that is based on the Minimum Description Length (MDL) principle. In a nutshell, it identifies as the best binning the one that best balances the complexity of the histogram and the likelihood of the data under this binning. In practice this means it automatically chooses both the number of and locations for the cut points that define the histogram depending on the complexity and size of the data.

Definition 2. **Representing Local Feature Distributions.** We first create a histogram for a given feature f_j and a set of nodes V with their feature values. A histogram consists a set of bins $b \in B_j$ each of which has a probability value based on the number of nodes corresponding to. Although we have chosen MDL binning to construct our histograms, FACETS will work with most histograms and binning approaches.

The **neighborhood (or local) distribution** $L_{i,j}$ is a distribution of features f_j over a set of neighbors of a particular node v_i; The **global distribution** G_j is the feature distribution across all nodes; and the **user profile distribution** U_j is the distribution of a sequence of interesting nodes V_h collected from the user during interaction with FACETS.

FACETS works by guiding users during their graph exploration using both surprisingness (Section 2.3) and subjective interest that changes dynamically to suit the user (Section 2.4). We do this comparing the local or neighborhood feature distributions with the global to determine surprisingness and the local with a user profile to determine dynamic subjective interest.

2.3 Ranking by Surprise

In order to calculate a node's surprisingness we compare the distribution of the node's neighbors with the global distribution for each feature. We chose a combined feature-centric and structural approach, because both structure and features play a critical role in inference problems [16]. Nodes whose local neighborhood vary greatly from the global are likely to be more surprising as they do not follow the general global trends.

One approach is to use the base entropy over node features to detect anomalous nodes; however, this ends up biasing the ranking towards skewed distribution. Instead we measure the difference between two distributions for more consistent results. Through our experiments we have chosen Jensen-Shannon (JS) divergence, a symmetrical version of Kullback-Leibler divergence to construct our surprisingness metric. JS divergence works well, because the resulting output is in bits so the divergences of several features can be easily combined into a single score. We measure surprise by determining the divergence of feature distributions $L_{i,j}$ over a node's neighborhood V_a (1 hop), from the global distributions of features G (see Equation 3). From these scores we select the top-k most surprising nodes (Equation 4).

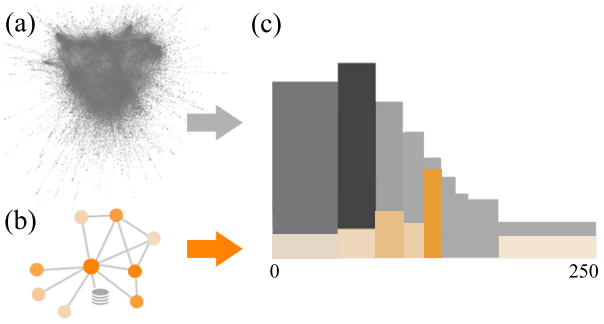


Figure 2: Local and global distribution histograms are both essential to FACETS. Local histograms (orange bars) are representations of feature distributions in a single node's egonet. The global distributions (gray bars) depicts the corresponding feature's distribution across the whole graph. The difference between those two distributions reflects if a node is "unusual" or surprising compared to the majority in the graph.

Given the JS Divergence or information radius between two distributions P and G:

$$D_{JS}(P||G) = \frac{1}{2}D(P||Q) + \frac{1}{2}D(G||Q), \tag{1}$$

where $Q = \frac{1}{2}(P + G)$ and $D(P||G)$ is the KL divergence for discrete distributions:

$$D(P||G) = \sum_b P(b) \log \frac{P(b)}{G(b)} \tag{2}$$

In Equation 2 we use base 2 so that $0 \leq D_{JS}(P||G) \leq 1$. For a fresh node v_a, whose neighbors are not yet visualized we first compute the surprise-score, s_i, of all neighboring nodes $v_i \in N(v_a)$:

$$s_i = \sum_{f_j \in A} \lambda_j D_{JS}(L_{i,j}||G_j), \tag{3}$$

where L_j and G_j are the local and global distributions of node-feature f_j and λ_j is a feature weight. Weighted feature scores in Equation 3 are used to lessen the impact of noisy features and to allow the user to lessen the contribution of a feature manually. The s_i scores are composed into $\hat{S}_a$, which holds all the scores for the neighbors of initial node v_a. We can find the most surprising k-nodes by looking for the largest divergence from the global:

$$\arg\max_{1\ldots k} \hat{S}_a \tag{4}$$

This yields the top-k most surprising nodes among the neighbors of node v_a. Since both the local-neighborhood feature distributions and the global feature distributions are static, the surprise scores can be precomputed to improve real time performance. We precompute and store surprise in FACETS to improve performance.

2.4 Ranking by Subjective Interestingness

We track the user's behavior and record a user profile as they explore their data. Each clicked node offers valuable details into the types of nodes in which the user is interested. This forms distributions U_j for each feature f_j.

To rank the user's interest in the undisplayed neighbors of node v_a we follow a similar approach as Equation 3:

$$r_i = \sum_{f_j \in A} \lambda_j D_{JS}(L_{i,j}||U_j), \tag{5}$$

where U_j is the distribution of feature f_j from the user's recent node browsing. In this case we want the local distributions that match better the user's current profile; i.e. we want the smallest possible divergences:

$$\arg\min_{1\ldots k} \hat{R}_a \tag{6}$$

Since this suffers from the cold-start phenomenon, because a user will not have a profile until they have explored some nodes, our remedy is to simply start the suggestions with surprising and important nodes, until the user has investigated several nodes.

3. OUR INTEGRATED APPROACH

The visual aspects of FACETS with the algorithmic rankings described previously. In this section we cover the graph visualization and design of our approach.

3.1 Components

We have created an adaptive graph exploration tool, FACETS, for performing fast and intuitizve exploration of graph datasets. FACETS was designed especially for graph tasks that require node-level details.

FACETS's user interface as shown in Figure 3 has four key elements: The first main area is the **Table View** (1), a conventional table showing the currently displayed nodes and some of their features. This provides sortable node-level information. The central area is the **Graph View** (2). It is an interactive *force-directed graph layout* that demonstrates the structure and relationships among nodes as the user explores. Coloring the nodes is used to encode the surprise and interest based on the user's current exploration. Visual scalability is an enormous challenge for nodes with many neighbors. Because we can only show a subset of the total nodes, we created the **Neighborhood Summary** (3). The neighborhood summary allows a user to investigate the feature distributions of its currently undisplayed neighbors. It presents the user with feature *heat maps*[2] that summarize the distributions of hidden nodes. When clicked, the heat maps turn into conventional distribution plots (histograms), where a user can compare the local neighborhood (orange) and the global (gray). This lets a user quickly select new nodes based on their feature values and get a quick summary of this node's neighborhood. As a user explores, we construct and display a summary profile of the important features they have covered in the **User Profile** view (4). The user profile view suggests high-level browsing behavior to the user; allows for better understanding of where the user-interest ranking comes from; and allows them to adjust if they want to ignore certain features in the interest ranking.

3.2 Design Rationale

In the following paragraphs, we discuss FACETS's contributions to graph exploration.

Exploring and Navigating

We use the term *graph navigation* to refer to the act of traversing graph data with a known destination or objective. *Graph exploration* is similar to navigation except that the user has little or no particular destination. One of our central design goals is to facilitate both exploration and navigation. We facilitate navigation through adaptation and exploration by filtering out unsurprising and unimportant nodes while still providing crucial feature details for hidden nodes via the Neighborhood Summary window. As shown in figure 5, the user can have a summarized view of a mouse-hovered nodes where the top ranked hidden neighbors, local distribution and global distribution are displayed. This lets the user be able to go through many nodes quickly in seek of nodes with desired feature and add those desired nodes to graph view easily from the top hidden node list or bin.

[2]While histogram encodes value of each bin as height, heat map uses darkness to represent values with equal height. The main advantage of heat map over histogram is its compact representation, which helps us save space.

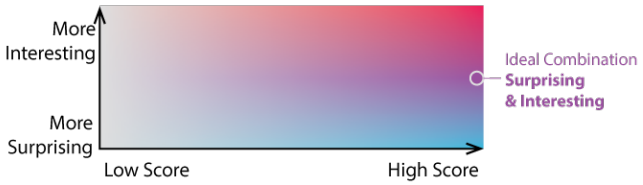


Figure 4: FACETS's color encoding of user-interest and surprise using hue and saturation. The horizontal axis indicates the overall score of a node; nodes with high scores will be located at the right of the chart and be colored with vibrant color. The vertical axis shows how much more user-interest score is than surprise score for a node.

Show the Best First

Keeping the graph view from becoming a hairball means only showing relevant, surprising, and interesting nodes. Importance, surprise, and user-interest are all important aspects of discovery, so we blend them into the results that are shown first to the user. Figure 4 illustrates how we visually encode the interest-surprise difference by hue and the sum of both scores by saturation. Nodes ranked high tend to have brighter color closer to purple, which becomes a clear visual cue for the user to quickly identify desired nodes. FACETS is almost completely free of parameters, making it simpler for users to explore their graphs.

Adaptive and Adjustable

Because user-interest varies greatly across users and even time, our design must be able to track the user's exploration behavior in order to approximate what is motivating them. Adapting as the user explores helps provide critical insight into users' latent objectives, because they can see how they have explored and also may find what they seek. During exploration, the users profile updates dynamically to illustrate a summary of their feature traversal, while the graph view provides the topological traversal. Even though there is no such necessarity to preset any parameters in order for our adaptive algorithm to work and the rankings are done in a black-box fashion during users' exploration, we allow them to directly manipulate the balance of features used in the interest calculation as well as the choice of which aspects form the ranking. Thereby when deciding to focus on only some of the features or weigh some feature less after having done some exploration with an established user profile, the user has the option to do so.

3.3 The FACETS Algorithm

In this section, we summarize the process of finding top-k most interesting and surprising neighbors in Algorithm 1. Whenever a user selects a node to explore, we rank its neighbors based on surprise and subjective interestingness we explained in Sec. 2.3 and 2.4. For each of the neighbors, we compute surprise and interest scores for each feature and aggregate them based on feature weights λ_j. We blend those scores, and finally, nodes with k highest scores will be returned.

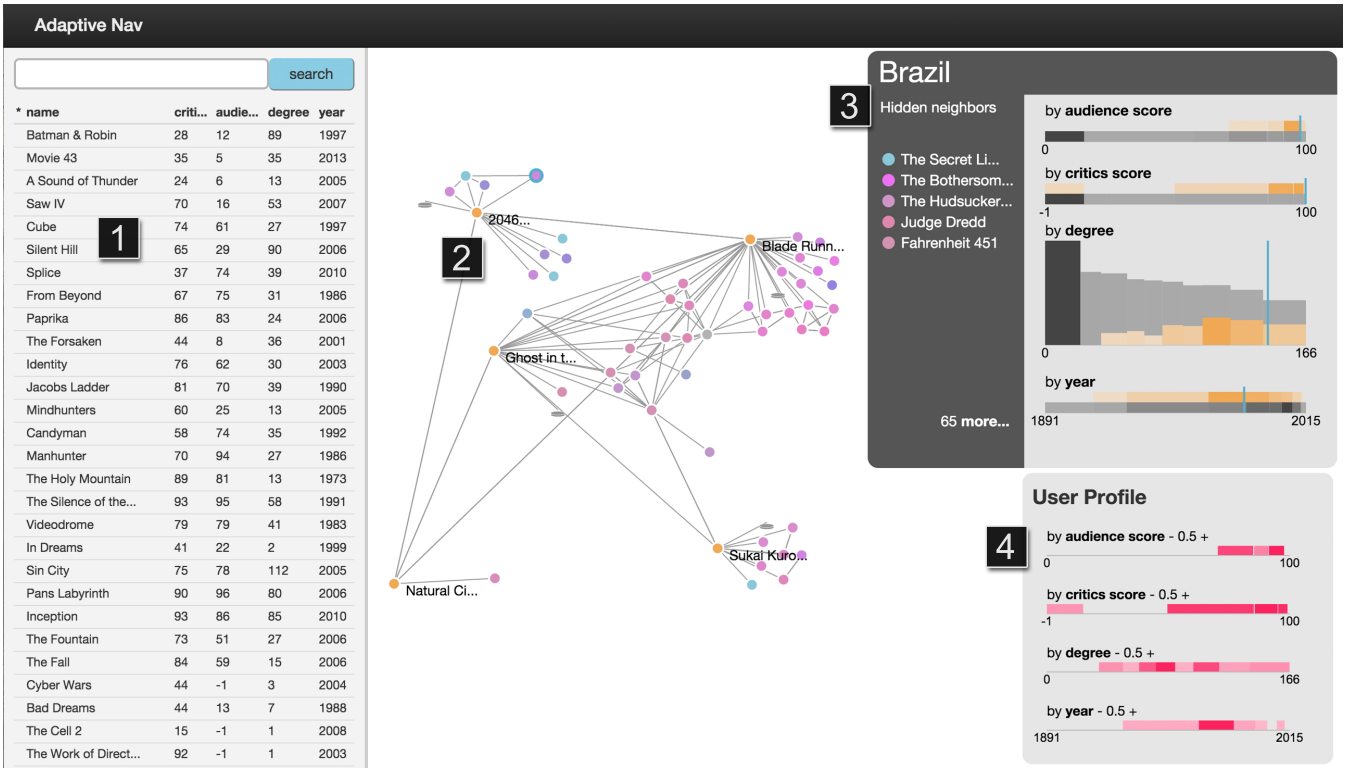


Figure 3: The FACETS user interface displaying an explored portion of the RottenTomatoes similar-movie graph. The user has traversed several films (shown as orange nodes with titles) and FACETS has displayed a subset of the relevant neighbors (the nodes with colors ranging from blue to red) and their connectivity. 1: The *Table View* provides a conventional summarization of the already explored nodes and some of their features. 2: The *Graph View* shows the connectivity of similar films as the user explores (it is linked with the table so that a selected node is highlighted in both views. 3: The *Neighborhood Summary* shows the currently hidden nodes, by their feature distributions, for a selected film. 4: The *User Profile* demonstrates a heat-map or flat histogram of the features the user has covered so far in their exploration. Figures are viewed best in color.

4. EXPERIMENTS

In this section, we evaluate the effectiveness and speed of FACETS. Our goal with FACETS is to facilitate user discovery by directing users towards regions of the graph with features suited to their interests. Our goal is *exploratory search* [18]; we are not trying to improve the accuracy of individual node recommendations, so it is not relevant to use canonical evaluation metrics like precision, recall, MAE, and RMSE [9, 29, 12]. Instead, we demonstrate the effectiveness of our approach with (1) three case studies that investigate the results of our algorithm and (2) a comparison of our scoring with canonical node ranking techniques.

4.1 Datasets

We consider four large graphs to evaluate our method. First, we use the RottenTomatoes[3] (RT) movie dataset, an attributed graph that contains basic information per movie (e.g., released year), as well as users' average ratings and critics scores. Second, we use the Google Web network, the DBLP co-authorship graph, and the YouTube network datasets from the SNAP repository[17]. We give the base characteristics of these graphs in Table 2. FACETS ignores nodes with zero degree, because they aren't very useful for graph exploration.

[3]http://rottentomatoes.com/

4.2 Runtime Analysis

Next, we evaluate the scalability of FACETS over several graphs. In particular, we will demonstrate its linear cost in the exploratory rankings. We note that the goal of this evaluation is mostly proof-of-concept – in this paper it is not our focus to get the lowest possible ranking time, which could be achieved via better hardware and further engineering.

Guided graph exploration requires sub-second rankings in order to remain smooth and reactive to user input. We must have rankings ready to display in under a second and ideally under half a second. This is why we have chosen to treat nodes in the tail of the degree distribution separately than their modest degree neighbors.

We have analyzed the runtime of FACETS, in Figure 6, using the graphs from Table 2; all but the RT graph used eight synthetic features. In our tests we use both random ordering and contiguous node ordering, displayed as Rand and Hop in Figure 6. Random ordering simulates using the search functionality while the hop ordering simulates hopping from one node to its neighbors during exploration. During the hop tests, the choice for the next hop must be among connected nodes that have not yet been traversed. High degree nodes have a higher chance of being selected and account for the fact that hop sometimes is slower than random in Figure 6. The graphs we tested demonstrate that the cost of the ranking is linear in the number of neighbors in the neighborhood. Our ranking requires both a value lookup and a single JS divergence calculation for each node and for each feature.

Algorithm 1: RankNeighbors

input : node v_a, precomputed histograms for: global feature distributions G and user profile distributions U
output: k most interesting and surprising neighbors of node v_a

if *deg(v_a) $\geq$ 1000* **then**
 $V_a \leftarrow$ top 1000 neighbors of v_a by highest degree
else
 $V_a \leftarrow$ all neighbors of v_a

forall the *nodes* v_i *in* V_a **do**
 forall the *features* f_j *in* A **do**
 $s_i^{(j)} = D_{JS}(L_{i,j}||G_j)$ // see Eq. 3
 $r_i^{(j)} = D_{JS}(L_{i,j}||U_j)$ // see Eq. 5
 $t_i^{(j)} = w_s s_i^{(j)} + w_r(1 - r_i^{(j)})$
 $\hat{T}_a[i] = \sum_{f_j \in A} \lambda_j t_i^{(j)}$
$T_{Nodes} = \arg\max_{1\ldots k} \hat{T}_a$

Network	Nodes	Edges
Rotten Tomatoes	17 074	72 140
DBLP	317 080	1 049 866
Google Web	875 713	5 105 039
Youtube	1 134 890	2 987 624

Table 2: Graph datasets used in our speed testing. They were picked for their variety in size and domain.

As mentioned in Section 2, the surprise scores are precomputed and can be accessed quickly. The cost to rank neighbors comes largely from the interest scores which cannot be precomputed. The cost, in JS divergence calculations, is $O(n \cdot f)$, and is asymptotically linear in both the number of neighbors n and the number of features f. Given our use of MDL histograms for the features, we are able to scale the number of features at low linear incremental cost (see Figure 7). Each neighbor only requires exactly one JS divergence calculation per feature (comparing the user profile and the local distribution).

Since many graphs contain triangles, it is very likely that redundant calls will be made during a user's exploration. We use this to our advantage and *cache* the distributions for each visited node rather than refetching them each time. Graphs with higher clustering coefficient may achieve better caching performance. For all but the Youtube graph, the caching became memoizing as the entirety of the nodes could fit in the cache.

4.3 Case Studies

Here, we describe three case studies using two graph datasets — the first two using the Rotten Tomatoes (RT) movie graph, and the last one using the DBLP co-authorship graph. These case studies illustrate how FACETS helps the user explore and graphs, incrementally gain understanding, discovers new insights, and visualize them. FACETS adapts to the user to provide surprising and interesting rankings to help the user explore relevant parts of the graphs.

Movie Example I: Blade Runner

Our first user John likes mid-90's post-apocalyptic and action films (shown as gray circles in Figure 8(a), e.g., *The Crow*, *Waterworld*), many of which were well received by critics and audience alike. John can add more films that he likes by exploring node-to-node or by the search utility. After exploring a few films in such genre, John is particularly interested in **Blade Runner**. Based on why

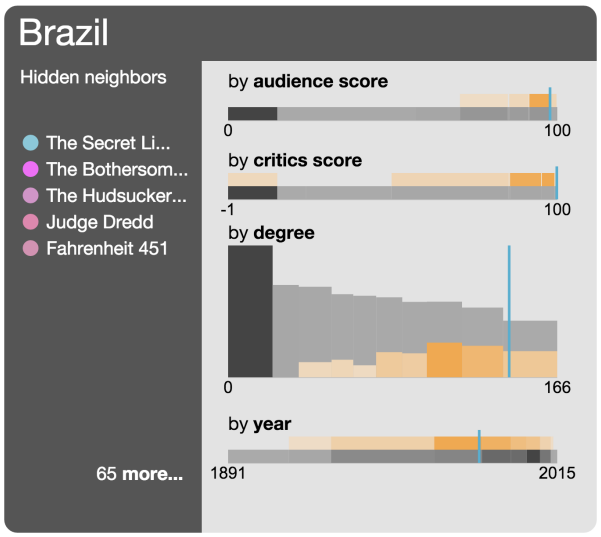


Figure 5: FACETS's neighborhood summary view, which displays the top ranked neighbors of the given node (left) and distributions of the current neighborhood's features (right). Each feature is displayed by a compact heat map, which can be expanded into a histogram. Each heat map shows both the global distribution (gray) and the node's neighborhood's distribution (orange).

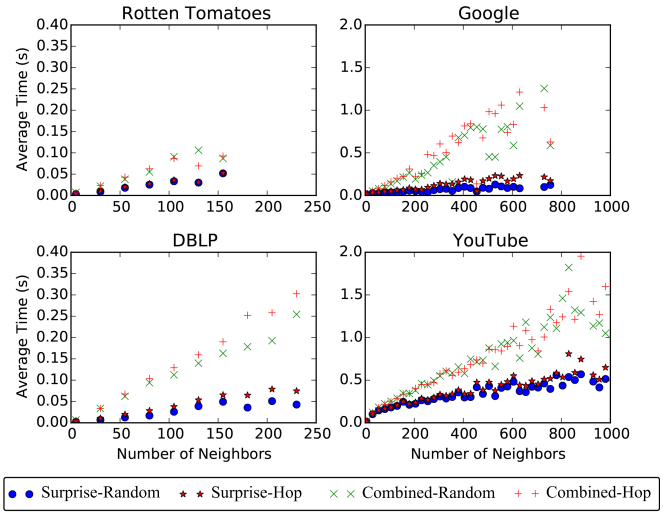


Figure 6: FACETS ranks neighbors in linear time, which is necessary to handle the large numbers of nodes that a user may explore. We show the average time to calculate the JS divergence for surprise and the combination of surprise and interest over a neighborhood of size n. FACETS uses the combined ranking if the user has explored enough to have a user profile from which we can extract the subjective interest ranking. We tested with contiguous node ordering to simulate normal exploration and random ordering to simulate a user searching using the table view.

	Example I					Example II				
Current Movies	**Blade Runner**					**Toy Story**				
Movies Visited	Waterworld, Braveheart, Pulp Fiction, The Crow, Fargo					A Bug's Life, Kung Pu Panda, Jumanji, The Incredibles, How to Train Your Dragon				
Top-k by	**Title**	**I**	**PR**	**BC**	**EV**	**Title**	**I**	**PR**	**BC**	**EV**
Interest	L.A. Confidential	1.06	2.23	264k	7.10	Monsters University	0.54	0.76	1070	29.3
	Sin City	1.15	4.34	833k	44.7	Rio	0.63	1.66	25464	57.4
	Dredd	1.20	1.74	55k	43.3	Toy Story II	0.70	1.84	55486	59.1
	V for Vendetta	1.23	2.99	620k	43.5	The Iron Giant	0.71	1.43	54079	46.3
	Heat	1.23	4.18	385k	11.0	ParaNorman	0.71	0.94	23291	15.1
Top-k by	**Title**	**S**	**PR**	**BC**	**EV**	**Title**	**S**	**PR**	**BC**	**EV**
Surprise	The Creation of the Humanoids	2.93	0.19	851	0.58	Buzz Lightyear of Star Command	3.15	0.14	12	2.14
	Demon Seed	2.75	0.50	650	1.54	Small Fry	3.01	0.17	132	3.10
	Natural City	2.69	0.28	437	0.74	Monsters University	2.64	0.76	1070	29.3
	Virtuosity	2.64	0.35	668	1.61	Cloudy With a Chance of Meatballs	2.46	1.28	1027	42.5
	Soylent Green	2.59	0.57	6578	1.50	ParaNorman	2.45	0.94	23291	15.1

Table 3: Comparing FACETS's surprise and interest ranking with common importance rankings (I&S is interest or surprise, PR is PageRank ($\times 10^{-4}$), BC is betweenness centrality, EV is eigenvector centrality ($\times 10^{-3}$). Each example has a selected film, a user profile at the time of selection, and the interesting and surprising neighbors.

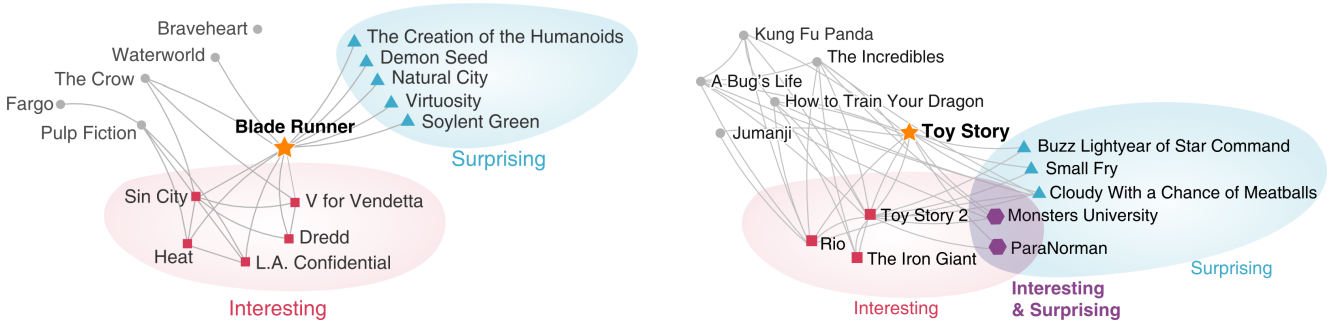


Figure 8: Visualizations of the (a) Blade Runner and (b) Toy Story case studies. The gray, circular nodes form a profile of films investigated by the user via graph exploration or text-search. The star nodes are the last-clicked node for which FACETS has produced the 5 most interesting (red), and surprising (blue) nodes. Nodes can be both surprising and interesting (purple).

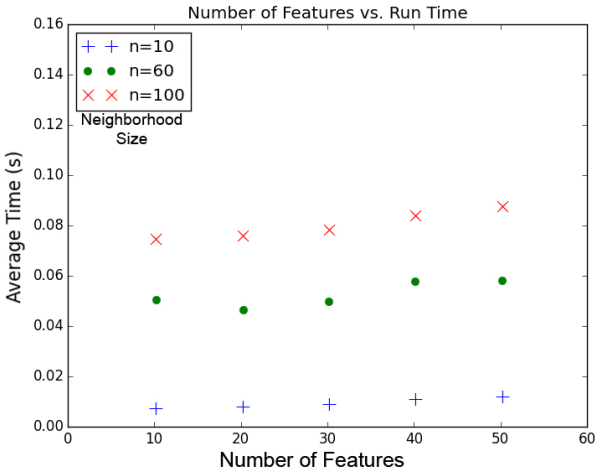


Figure 7: FACETS scales linearly in the number of features, because the ranking requires exactly one JS divergence calculation per neighbor for interest and constant time lookup for surprise. Because of the compact structure of our MDL bins, the incremental cost of adding additional features is quite low.

John has explored, FACETS returns top-ranking surprising and interesting films. The top-5 of each kind are displayed in Table 3 on the left and in Figure 8(a). These movies' interest (I) and surprisingness (S) scores are determined based on conventional measures of node importance; PageRank (PR), betweenness centrality (BC), and eigenvector centrality (EV).

Many of the surprising movies would not considered important by canonical approaches, partly because FACETS's surprise rank operates on both features and local graph structure rather than global structure. Movies that are very heavily connected also face a higher chance of matching the global distribution and therefore being less surprising.

While the notion of surprise is often considered "unimportant" by conventional metrics, interest exhibits more variety in importance than surprise. This is especially apparent in both this and the next case studies. Here we have selected types of films that have really large viewership and represent a very large genre in modern film with many potentially similar movies. Both the structure and features used in our subjective interest have lead the ranking towards more conventionally important nodes (consider the high PR, BC, and EV scores).

As the highest possible JS-divergence for a single feature is 1; the maximum divergence is the sum of feature weights Λ. In this case study, we used five features, so interest scores ≤ 1.0 suggest

that the user profile is in relatively good agreement with the proposed node. Nodes with very low interest divergence are strong candidates that their neighborhood will be of interest to the user.

Movie Example II: Toy Story

Our second user Bonnie investigates several children's films that were criticized by the critics, but still enjoyed by audiences (all have a lower critics' score than audience score). Bonnie hops from node-to-node across the listed films in the order they appear in Table 3. She then selects **Toy Story**. The results are displayed in Table 3 on the right and in Figure 8(b).

As in the previous example, the surprising nodes tend not to be conventionally important. Yet, consider the difference in importance of the interesting films versus the previous example, many of these suggestions have significantly lower importance. The second profile has less coherent features and they do not draw the the interest ranking towards conventionally important nodes, despite that *Toy Story* is a very well connected node. Also note that **Monsters University** and **ParaNorman** are featured in both the top-5 surprising and interesting movies. The surprise score and interest score are not counter to each other. In fact, this is an excellent example, because these films are both surprising and subjectively interesting!

DBLP Example III

Our third example uses data extracted from DBLP, a computer science bibliography website [4]. The graph is an undirected, unweighted graph describing academic co-authorship. Nodes are authors. An edge connects two authors who have co-authored at least one paper.

Our user Jane is a first-year graduate student new to data mining research. She just started reading seminal articles written by Philip Yu (topmost orange star in Figure 9). FACETS quickly helps Jane identify other prolific authors in the data mining and database communities, like *Jiawei Han*, *Rakesh Agrawal*, *Raghu Ramakrishnan*, and *Christos Faloutsos*; these authors have similar feature distributions as Philip Yu (e.g., very high degree). Jane chooses to further explore *Christos Faloutsos*'s co-authors. FACETS suggests *Duen Horng Chau* as one of the surprising co-authors, who seems to have relatively low degree (i.e., few publications) but has published with highly-prolific co-authors. Among these is *Brad Myers* (leftmost orange start in Figure 9), who publishes not in data mining, but in human-computer interaction (HCI). This exploration introduces Jane to a new field, and she wants to learn more. Using FACETS's interest-based suggestion, she discovers a community of co-authors who have published with Brad; among them, *Mary Beth Rosson* further leads to another community of HCI researchers, which includes *Ben Shneiderman*, the visualization guru!

5. RELATED WORK

Graph traversal and exploration has been investigated in several different areas related to networks.

Trails and Paths Through Networks

The information retrieval community has proven an asset in analyzing the web-browsing paths users covered as they explored the web. *Click trail analysis* has been used to analyze the website-to-website paths of millions of users in order to improve the ranking of search results [4, 24].

Intermediate sites and destinations of common trails can be used successfully as search results themselves as in [31]. If a user is following a common trajectory they can even be "teleported" to the destination page [26]. However, in our case we do not have millions of explored paths through our input network and cannot directly rely on the aggregate analysis of trails used above.

West et al. analyzed Wikipedia users' abilities and common patterns as they explored Wikipedia [30]. They observed that users would balance between a conceptually simple solution at the cost of efficiency–there were more direct routes to their target article that made less sense than they path they chose. Their system also used trail analysis in order to try to predict where a user would go based on the user's article-trail features.

[4]http://dblp.uni-trier.de

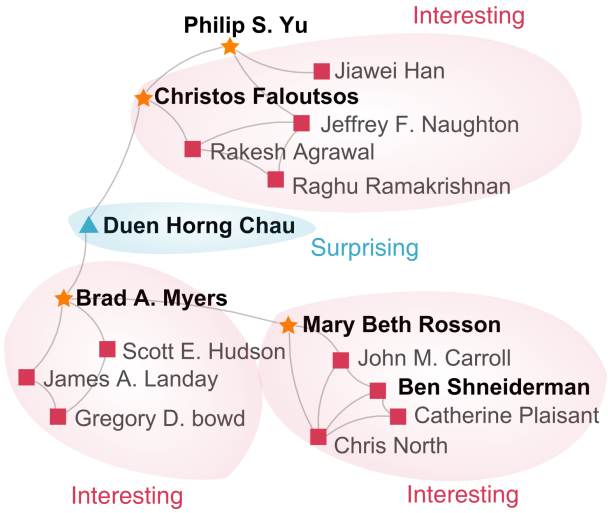


Figure 9: Visualization of our user Jane's exploration of the DBLP co-authorship graph. Jane starts with Philip Yu. FACETS then suggests Christos Faloutsos and several others as prolific data mining researchers. Through Christos, FACETS suggests Duen Horng Chau as a surprising author as he has published with both data mining and human-computer interaction (HCI) researchers, like Brad Myers. Through Brad, FACETS helps Jane discover communities of HCI researchers, including Ben Shneiderman, the visualization guru.

Degree of Interest

The visualization community has also investigated local graph exploration. Bottom-up exploration first appeared in [10], a tool for exploring hierarchies using a "degree of interest" (DOI) function to rank the relevance of currently undisplayed nodes. The idea of DOI was later expanded by [28] to apply to a greater set of graph features. The Apolo system [5] further improves on it to allow users to freely define their own arbitrary number of clusters, which it uses to determine what to show next, through the Belief Propagation algorithm. Recently, the DOI idea is applied to time-varying settings, to capture salient temporal graph changes [1].

We have built on the idea of using a DOI to determine the ranking for what nodes we show users; however, we use a *dynamic DOI* function which changes to suit the browsing behavior of the users as they explore their data.

Surprise and Serendipity

Algorithms like Oddball [2], an unsupervised approach to detect anomalies in weighted graphs, can be used to detect surprising nodes. Akoglu et al. focused on anomalies based on edge weight-

ing and not on node-level features. The TANGENT algorithm by Onuma et al. is a parameter free technique used to discover surprising recommendations [19]. They measure the surprise in their model by measuring the amount of horizon-broadening each new node imparts, where the horizons are edges going to different clusters of users and films.

Similar to our approach to comparing local and global distributions, Parameswaran et al. [21] presented the idea of finding interesting dimensions in the context of data cube by comparing their distributions to those aggregated.

Andre et al. [3] studied the potential for using serendipity in Web search and the effects of personalization on that potential. They found out that serendipity could be useful for certain queries, and personalization, ideas which we leverage in FACETS.

De Bie [6, 7] proposed a general framework for measuring the interestingness of data mining results as their log-likelihood given a Maximum Entropy (MaxEnt) distribution based on the background knowledge of the user. Instead, we consider the Jensen-Shannon divergence between the local and global histograms as the surprisingness of a node. Our notion of subjective interestingness comes closer to that of Tatti and Vreeken [25], whom proposed to iteratively update the model, to avoid redundancy. Our goal is orthogonal, as we are specifically interested in identifying nodes that are similar to those the user chose to explore.

6. CONCLUSION

In this work, we presented FACETS, an integrated approach that combines visualization and computational techniques to help the user performs adaptive exploration of large graphs. FACETS overcomes many issues commonly encountered when visualizing large graphs. FACETS shows the user only the most subjectively interesting material as they explore. We do this by ranking the neighbors of each node by surprisingness and subjective interest based on what the user has explored so far. The surprisingness rankings are determined by the divergence of local feature distributions from the background feature distributions. The subjective interest rankings arise from the divergence between the local feature distributions and the user's current profile, which represents their exploration, in node-features, up to that point.

Our FACETS algorithm is scalable and is linear in the number of neighbors and linear in the number of features. We demonstrated the effectiveness of FACETS through case studies using the RottenTomatoes movie graph and the DBLP co-authorship graph, and comparison with canonical importance ranking measures. The user interest ranking influences the direction of a user's exploration by showing the best matches to their current taste in nodes. Despite the old adage, you can see the graph through its nodes.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1217559 and National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1148903. Funding was provided by the U.S. Army Research Office (ARO) and Defense Advanced Research Projects Agency (DARPA) under Contract Number W911NF-11-C-0088. Jilles Vreeken is supported by the Cluster of Excellence "Multimodal Computing and Interaction" within the Excellence Initiative of the German Federal Government. James Abello acknowledges support from mgvis.com, DIMACS, and NSF grant CCF-1445755.

7. REFERENCES

[1] J. Abello, S. Hadlak, H. Schumann, and H.-J. Schulz. A modular degree-of-interest specification for the visual analysis of large dynamic networks. *Visualization and Computer Graphics, IEEE Transactions on*, 20(3):337–350, 2014.

[2] L. Akoglu, M. McGlohon, and C. Faloutsos. Oddball: Spotting anomalies in weighted graphs. In *PAKDD*, pages 410–421. Springer, 2010.

[3] P. André, J. Teevan, and S. T. Dumais. From x-rays to silly putty via uranus: serendipity and its role in web search. In *SIGCHI*, pages 2033–2036. ACM, 2009.

[4] M. Bilenko and R. W. White. Mining the search trails of surfing crowds: Identifying relevant websites from user activity. In *WWW*, pages 51–60. ACM, 2008.

[5] D. H. Chau, A. Kittur, J. I. Hong, and C. Faloutsos. Apolo: interactive large graph sensemaking by combining machine learning and visualization. In *KDD*, pages 739–742. ACM, 2011.

[6] T. De Bie. An information theoretic framework for data mining. In *KDD*, pages 564–572. ACM, 2011.

[7] T. De Bie. Maximum entropy models and subjective interestingness: an application to tiles in binary databases. *Data Min. Knowl. Disc.*, 23(3):407–446, 2011.

[8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, volume 29, pages 251–262. ACM, 1999.

[9] F. Fouss, A. Pirotte, J.-M. Renders, and M. Saerens. Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation. *IEEE TKDE*, 19(3):355–369, March 2007.

[10] G. W. Furnas. Generalized fisheye views. *SIGCHI Bull.*, 17(4):16–23, Apr. 1986.

[11] T. H. Haveliwala. Topic-sensitive pagerank. In *WWW*, pages 517–526, New York, NY, USA, 2002. ACM.

[12] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *TOIS*, 22(1):5–53, 2004.

[13] D. A. Keim. Visual exploration of large data sets. *Commun. ACM*, 44(8):38–44, 2001.

[14] D. A. Keim. Information visualization and visual data mining. *IEEE TVCG*, 8(1):1–8, 2002.

[15] P. Kontkanen and P. Myllymäki. MDL histogram density estimation. In *AISTATS*, 2007.

[16] T. Lee, Z. Wang, H. Wang, and S. won Hwang. Attribute extraction and scoring: A probabilistic approach. In *ICDE*, pages 194–205, 2013.

[17] J. Leskovec and A. Krevl. SNAP Datasets: Stanford large network dataset collection. `http://snap.stanford.edu/data`, June 2014.

[18] G. Marchionini. Exploratory search: from finding to understanding. *Commun. ACM*, 49(4):41–46, 2006.

[19] K. Onuma, H. Tong, and C. Faloutsos. Tangent: a novel,'surprise me', recommendation algorithm. In *KDD*, pages 657–666. ACM, 2009.

[20] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bringing order to the web. Technical report SIDL-WP-1999-0120, Stanford, 1999.

[21] A. Parameswaran, N. Polyzotis, and H. Garcia-Molina. Seedb: Visualizing database queries efficiently. *PVLDB*, 7(4), 2013.

[22] C. Plaisant, B. Milash, A. Rose, S. Widoff, and B. Shneiderman. Lifelines: visualizing personal histories. In *SIGCHI*, pages 221–227. ACM, 1996.

[23] B. Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *VL*, pages 336–343. IEEE, 1996.

[24] A. Singla, R. White, and J. Huang. Studying trailfinding algorithms for enhanced web search. In *SIGIR*, pages 443–450. ACM, 2010.

[25] N. Tatti and J. Vreeken. Comparing apples and oranges - measuring differences between exploratory data mining results. *Data Min. Knowl. Disc.*, 25(2):173–207, 2012.

[26] J. Teevan, C. Alvarado, M. S. Ackerman, and D. R. Karger. The perfect search engine is not enough: A study of orienteering behavior in directed search. In *SIGCHI*, pages 415–422. ACM, 2004.

[27] H. Tong, C. Faloutsos, and J.-Y. Pan. Fast random walk with restart and its applications. In *ICDM*, pages 613–622. IEEE, 2006.

[28] F. Van Ham and A. Perer. Search, show context, expand on demand : Supporting large graph exploration with degree-of-interest. *IEEE TVCG*, 15(6):953–960, 2009.

[29] F. Wang, S. Ma, L. Yang, and T. Li. Recommendation on item graphs. In *ICDM*, pages 1119–1123, Dec 2006.

[30] R. West and J. Leskovec. Human wayfinding in information networks. In *WWW*, pages 619–628. ACM, 2012.

[31] R. W. White and J. Huang. Assessing the scenic route: Measuring the value of search trails in web logs. In *SIGIR*, pages 587–594. ACM, 2010.